

Lecture 3: Aggregate Production, and Cross-Country Differences in L , K , and A

Christian Wolf

14.05, Fall 2025

Today's lecture

- Today we begin with the growth half of the course
- Long-run objective: understand “**deep**” **determinants of income**
 - Why are some countries rich, some countries poor?
 - Why do countries grow over time?
- Today we begin “**proximate**” **causes of income**: relate income to factors of production
 - In the Robinson-Crusoe economy income (= coconuts) was a function of hours worked and productivity. Today we'll be more general: add capital (= machines).
 - Only job today is to understand this production function. We'll later embed it in a proper general equilibrium macro model with interesting labor supply/savings decisions.

Today's lecture

1. Introduce the concept of an **aggregate production function**

- Express income as a function of a small number of production factors (labor, capital) + “productivity”
- Aside: does it make sense to aggregate so much? there are many different kinds of labor, many different machines, ...

2. Discuss what **shape** of that production function makes empirical sense

3. Decompose cross-country differences in income into **differences in factors** (capital per person vs. “productivity”)

- Takeaway will be: poor countries are poor because they have fewer/worse machines and use them less well
- Will lead us to next **Q**: why do they have fewer/worse machines and use them less well?

Outline

1. A neoclassical production model

The aggregate production function

Firm behavior

Household behavior

Equilibrium

2. Aside: A single good?

Planner's problem vs. competitive equilibrium

Empirical assessment

3. Using the production model to interpret cross-country income differences

A stylized model of production

- We relate **output/income** to factors of production:
 1. **labor** (L)
 2. **capital** (K)
 3. a residual that captures **productivity/technology** (A)
- This certainly makes sense for a single good, like the coconuts in the Robinson Crusoe economy. So let's for now consider a one-good economy. Later: can we generalize?
 - There are a fixed number of people available to produce, $\bar{L}$. People also have a fixed number of machines, $\bar{K}$.
 - Firms hire people and machines to produce the final good.
- **Q:** How much of the good gets produced? How much do capital and labor get paid?
A via 4-step approach: technology, firm & household behavior, "equilibrium". Just like what we did for the competitive Robinson Crusoe economy.

Outline

1. A neoclassical production model

The aggregate production function

Firm behavior

Household behavior

Equilibrium

2. Aside: A single good?

Planner's problem vs. competitive equilibrium

Empirical assessment

3. Using the production model to interpret cross-country income differences

Aggregate production function

- The technology for producing the single good is given by:

$$Y = A \cdot F(K, L). \quad (1)$$

- Here, $F : \mathbb{R}_+^2 \rightarrow \mathbb{R}_+$ is an **aggregate production function**.
 - Specified for all possible K, L , not just $\bar{K}, \bar{L}$ that is available.
 - A scales how productively the factors are used—“total factor productivity”. Ignore for now.
- We assume that F is continuous and twice differentiable, and we let $F_x \equiv \partial F / \partial x$, $F_{xx} \equiv \partial^2 F / (\partial x \partial x)$ denote the derivatives.
- We will also assume that the production function F is “neoclassical” in the sense defined on the next slide

A “neoclassical” production function

Definition

The production function, F , is **neoclassical** if it satisfies the following:

1. Constant returns to scale, or linear homogeneity:

$$F(\lambda K, \lambda L) = \lambda F(K, L), \quad \forall \lambda \geq 0.$$

2. Positive but diminishing marginal products:

$$\begin{aligned} F_K(K, L) &> 0, & F_L(K, L) &> 0, \\ F_{KK}(K, L) &< 0, & F_{LL}(K, L) &< 0. \end{aligned}$$

Here, $F_x \equiv \partial F / \partial x$, $F_{xx} \equiv \partial^2 F / (\partial x \partial x)$ denote the derivatives.

3. Inada conditions:

$$\lim_{K \rightarrow 0} F_K = \lim_{L \rightarrow 0} F_L = \infty \quad \lim_{K \rightarrow \infty} F_K = \lim_{L \rightarrow \infty} F_L = 0.$$

Constant returns to scale

- Constant returns to scale (CRS) makes our life easy for two reasons:
 1. It supports perfect competition: when we discuss firm behavior, price-taking (= many small firms) will make sense. With increasing returns, a single firm would take over.
 2. It allows us to reduce dimensionality by dividing all variables by L . We can thus normalize size, facilitating cross-country comparisons. See below.
- Define output and capital per labor as $y \equiv Y/L$ and $k \equiv K/L$. We can then define the **intensive (or per labor) production function** as:

$$f(k) \equiv F(k, 1).$$


By CRS,

$$Y = F(K, L) = L \cdot F(K/L, 1) = L \cdot f(k)$$

- This verifies that aggregate production is equivalently described by:

$$\underbrace{y}_{\text{goods per labor}} = f \left(\underbrace{k}_{\text{capital per labor}} \right)$$

Illustration



Example: Cobb-Douglas production function

$$F(K, L) = K^\alpha L^{1-\alpha}$$

- Here, $\alpha \in (0, 1)$ is output's elasticity with respect to capital.
- Check that $F(K, L)$ satisfies all the assumptions so far. Also check that the intensive form function is given by:

$$f(k) = k^\alpha.$$

- The Cobb-Douglas functional form useful because of its tractability and also because it is a good proxy for reality—more on this shortly.
- But let us first go back to our production model and let us understand how firms behave given some production function F .

Outline

1. A neoclassical production model

The aggregate production function

Firm behavior

Household behavior

Equilibrium

2. Aside: A single good?

Planner's problem vs. competitive equilibrium

Empirical assessment

3. Using the production model to interpret cross-country income differences

The representative firm

- Suppose that there are many firms, all with the same production function $F(K, L)$.
- Let the rental rates of capital and labor (wages) be R and w . All firms face those same prices. Also normalize the price of the output good to 1, as in the RC economy.
- We can represent this situation *as if* there is a *single* firm.
 - Why? all firms are identical, so they all make the same choices
 - But crucially: this hypothetical “single firm” is not a monopolist. It doesn’t set prices, it takes them as given. Makes sense because of CRS.
- This is the concept of the economy’s “representative” firm.

The representative firm

- The representative firm chooses the levels of its inputs and its production so as to maximize profits, taking R and w as given.
- That is, the firm's problem is:

$$\Pi = \max_{K,L} \{F(K, L) - RK - wL\}. \quad (2)$$

- Profits are distributed back to households (firms' owners), so the following accounting identity holds:

$$\underbrace{Y}_{\text{national income}} = \underbrace{RK}_{\text{capital's take}} + \underbrace{wL}_{\text{labor's take}} + \underbrace{\Pi}_{\text{profits}}.$$

- We are interested in understanding how Y is determined and how this is distributed among factors (i.e., labor and capital) and profits.


Firms demand capital and labor

- By the concavity of F , the following first-order conditions (FOCs) are both necessary and sufficient for a solution to the firm's problem:


$$F_K(K, L) = R \text{ and } F_L(K, L) = w.$$

- The first condition states that the marginal product of capital is equated with its rental rate. The second condition does the same for labor and the wage.
 - In short, the price of each input is equated with its marginal product.
 - This is like what we saw for Crusoe Inc., just now there are two inputs.
- If the price of an input exceeded its marginal product, the typical firm would have found it profitable cut down on its use; and if the converse were true, the typical firm would rather expand the use of this input.
- These equations describe declining demand curves ...

Firms demand capital and labor



Demand for capital



Demand for labor

- This describes one half of the market—firm behavior gives declining demand curves.
- Remains to specify the other half—labor and capital supply.

Outline

1. A neoclassical production model

The aggregate production function

Firm behavior

Household behavior

Equilibrium

2. Aside: A single good?

Planner's problem vs. competitive equilibrium

Empirical assessment

3. Using the production model to interpret cross-country income differences

Households supply capital and labor

- In our stylized model, households supply all of their labor and capital, $\bar{L}$ and $\bar{K}$, as long as $w > 0$ and $R > 0$.
- This is because they have no alternative option—a simplifying assumption that we will relax soon.
 - Of course we did already relax it in the Robinson Crusoe economy, where Robinson provided labor, trading off consumption and leisure, rather than just providing $\bar{L}$.
 - We are making our lives easier here because at the moment we just want to focus on the production function, nothing else.
- Combining the supply with demand, we obtain the equilibrium ...

Outline

1. A neoclassical production model

The aggregate production function

Firm behavior

Household behavior

Equilibrium

2. Aside: A single good?

Planner's problem vs. competitive equilibrium

Empirical assessment

3. Using the production model to interpret cross-country income differences

Equilibrium: definition

equilibrium = optimality + market-clearing


Exactly as in the Robinson Crusoe economy, just even easier here.

Definition


A competitive equilibrium is a collection of prices $\{R^*, w^*\}$ as well as allocations $\{K^*, L^*, Y^*, \Pi^*\}$ such that:

1. Households optimize, which here simply means that they supply $\bar{K}$ and $\bar{L}$.
2. Firms optimize, that is, (Π^*, K^*, L^*) solves (2) given $R = R^*$ and $w = w^*$.
3. Markets clear, that is, $K^* = \bar{K}$ and $L^* = \bar{L}$.

Equilibrium: graphical illustration



Equilibrium in the capital market



Equilibrium in the labor market

Equilibrium: basic implications

- In equilibrium, the economy ends up utilizing its factors:

$$Y^* = F(\bar{K}, \bar{L}) .$$

- And factors end up getting paid their marginal products:

$$R^* = F_K(\bar{K}, \bar{L}) \text{ and } w^* = F_L(\bar{K}, \bar{L}) . \quad (3)$$

- What does this analysis imply for equilibrium profits, Π^* ?

CRS + competition $\Rightarrow$ profits are zero

- Recall constant returns to scale, $F(\lambda K, \lambda L) = \lambda F(K, L)$. This implies:

$$F(K, L) = F_K(K, L)K + F_L(K, L)L$$

Proof: differentiate $F(\lambda K, \lambda L)$ w.r.t. λ and evaluate at $\lambda = 1$.

- Combining this with (3) implies:

$$\Pi = Y - F_K K - F_L L = 0$$

or equivalently

$$Y = R^* \bar{K} + w^* \bar{L}.$$

- Profits are zero, all income distributed between capital and labor.
 - Unrealistic but simple. Profits that firms earn in the real world are captured inside $R^* \bar{K}$.
 - Aside: is $(\bar{K}, \bar{L})$ the only optimum for firms given (w^*, R^*) ?

Equilibrium with Cobb-Douglas production

- Let's characterize the solution further in the Cobb-Douglas case:

$$Y = F(K, L) = AK^\alpha L^{1-\alpha}$$

- The rental rate of capital is:

$$R = F_K = \alpha \frac{Y}{K}$$

- So the share of capital (i.e., capital's take as fraction of Y) is:

$$\frac{RK}{Y} = \frac{F_K K}{Y} = \alpha$$

- Similarly, the share of labor (i.e., labor's take as fraction of Y):

$$\frac{wL}{Y} = \frac{F_L L}{Y} = 1 - \alpha$$


Equilibrium with Cobb-Douglas production

- So in the Cobb-Douglas case, we have the decomposition

$$\begin{aligned} Y^* &= R^* \bar{K} + w^* \bar{L}, \\ &= \alpha Y^* + (1 - \alpha) Y^*. \end{aligned}$$

- Factor shares of income are constant regardless of $\bar{K}, \bar{L}$.
 - The factor shares are captured by the coefficients α and $1 - \alpha$.
 - Intuitively, think of shares as reflecting relative “importance.”
- We actually observe factor shares in the data ...

Labor's share of income



- For our purposes: can think of labor share as roughly a constant, around 0.60 or $2/3$.
- But: there is a worrisome (?) secular trend. Labor share seems to be declining in the last few decades. Reversal after Covid?

Labor's share & α

- Labor's share is around $2/3$. Capital's share (broadly construed) is the residual, $1/3$.
- So if we use Cobb-Douglas, and assume perfect competition, we need

$$\alpha \simeq 1/3$$

to fit observed factor shares.

- Aside: what about imperfect competition and the fact that actual firm profits are often far from zero?
 - It gets tricky, because we need to decide whether profits represent returns to inputs (such as R&D and entrepreneurial activity) or pure rents.
 - It is also quite telling that the recent decline in the labor share appears to be correlated with an increase in market concentration and markups. But we will abstract from these issues.

Taking stock from the production model

- In equilibrium, the economy ends up utilizing its factors:

$$Y^* = F(\bar{K}, \bar{L})$$

This of course was obvious from the beginning. We know total labor and capital, and we were given the production function.

- Factors end up getting paid their marginal products:

$$R^* = F_K(\bar{K}, \bar{L}) \text{ and } w^* = F_L(\bar{K}, \bar{L})$$

This is what we actually got from our analysis—the prices that *support* the outcome Y^* as the equilibrium production level.

- All income is distributed between capital and labor:

$$Y^* = R^*\bar{K} + w^*\bar{L}$$

- With Cobb-Douglas, capital & labor shares are α and $1 - \alpha$. Factor shares in the data discipline the CD coefficient, $\alpha \simeq 1/3$.

Outline

1. A neoclassical production model

The aggregate production function

Firm behavior

Household behavior

Equilibrium

2. Aside: A single good?

Planner's problem vs. competitive equilibrium

Empirical assessment

3. Using the production model to interpret cross-country income differences

Aggregate production function

- To simplify the analysis, we have assumed that there is a single good in the economy.
GDP = physical quantity of that good.
- Technology for this good gives GDP as a function of factors of production.
- What does this function represent in real economies with many goods, different sectors, different technologies, and different factor mixes?
 - Will show: if economy is “efficient”, then this single good view is actually quite general.
 - At any point in time, the productive factors of the economy (capital and labor) are allocated efficiently to sectors.
 - The solution to this assignment problem is what we have called $F(\bullet)$ —a mapping from resources to total production.
 - What happens if an economy is not “efficient”? Through the lens of our single-good accounting approach, all deviations from efficiency will go into the TFP term A .

Outline

1. A neoclassical production model

The aggregate production function

Firm behavior

Household behavior

Equilibrium

2. Aside: A single good?

Planner's problem vs. competitive equilibrium

Empirical assessment

3. Using the production model to interpret cross-country income differences

A planner's problem: production efficiency

- Suppose there are multiple sectors, indexed by $j \in \{1, \dots, J\}$, each facing a price P_j and using a production function F_j .
- First consider the problem of a fictitious “planner” that allocates capital and labor across the sectors so as to maximize GDP. [This is like the “optimal allocation” with RC.]

$$\begin{array}{ccc} \max & \sum_j \frac{P_j}{P} Y_j & \text{subject to} \\ & \underbrace{Y_j = F_j(K_j, L_j)}_{\text{technology for sector } j} \quad \forall j, & \underbrace{\sum_j K_j \leq \bar{K}}_{\text{constraint for capital}}, \quad \text{and} \quad \underbrace{\sum_j L_j \leq \bar{L}}_{\text{constraint for labor}} \end{array}$$

where P is a price index, so that $\frac{P_j}{P}$ is the relative price of good j . Normalize $P = 1$. Don't worry about the prices. They would come from household preferences over the goods. Also here there is one type of labor & capital. The math would be similar with multiple types.

- This gives GDP as increasing function of $\bar{K}$ and $\bar{L}$ = “aggregate production function”

A planner's problem: production efficiency

- Let's first characterize the solution to the planning problem.
- Write the Lagrangian of the planner's problem as follows:

$$\mathcal{L} = \sum_j P_j F_j(K_j, L_j) + \lambda_K (\bar{K} - \sum_j K_j) + \lambda_L (\bar{L} - \sum_j L_j)$$

- The FOCs give, for every j ,

$$P_j \frac{\partial F_j(K_j, L_j)}{\partial K} = \lambda_K \quad \text{and} \quad P_j \frac{\partial F_j(K_j, L_j)}{\partial L} = \lambda_L$$

- The planner's solution is then given by a combination of Lagrange multipliers (λ_K, λ_L) and cross-sectoral allocations $(K_j, L_j)_{j=1}^J$ that solve the above FOCs along with the two resource constraints for the two factors:

$$\sum_j K_j = \bar{K} \quad \text{and} \quad \sum_j L_j = \bar{L}$$

Competitive equilibrium

- Now consider the competitive eq'm. In each sector j , firms maximize profits,

$$P_j F_j(K_j, L_j) - RK_j - wL_j$$

- The FOCs give

$$P_j \frac{\partial F_j(K_j, L_j)}{\partial K} = R \quad \text{and} \quad P_j \frac{\partial F_j(K_j, L_j)}{\partial L} = w$$

which simply means that the marginal revenue product of each factor is equated to its price in all sectors.

- An equilibrium is given by a combination of factor prices (R, w) and allocations $(K_j, L_j)_{j=1}^J$ that solve the above FOCs along with the market-clearing conditions

$$\sum_j K_j = \bar{K} \quad \text{and} \quad \sum_j L_j = \bar{L}$$

- These are the same as the planner's resource constraints, except that now they can be read as "demand equals supply."

Planning problem = competitive equilibrium

- Compare the system of equations characterizing the planner's solution with that for the competitive equilibrium. The two systems coincide, and so do their solutions.

Decentralization: prices = planner Lagrange multipliers.

- They both deliver the same cross-sectoral allocation of factors, and they have the property that the factor prices coincide with the planner's shadow values of the factors:

$$R = \lambda_K \quad \text{and} \quad w = \lambda_L$$

- Interpretation
 - “maximizing firm profits” = “maximizing GDP” = aggregate production function to turn K and L into overall output
 - Yet another example of the welfare theorems/invisible hand: equivalence between (i) decentralized competitive equilibria and (ii) centralized planning problems

Outline

1. A neoclassical production model

The aggregate production function

Firm behavior

Household behavior

Equilibrium

2. Aside: A single good?

Planner's problem vs. competitive equilibrium

Empirical assessment

3. Using the production model to interpret cross-country income differences

Production efficiency in reality

- Efficient allocation of labor and capital across sectors and firms is not realistic, for many reasons, reflecting a failure of the welfare theorems:
 - monopoly power
 - taxes and regulation (good and bad)
 - financial constraints
 - limited information or managerial competence
 - ...
- All these “frictions” mean that, in practice, marginal productions are not equalized, and so GDP is not maximized.
- Plus all of this is static. Maybe dynamically investment in the “wrong” kinds of capital goods, technologies, types of workers?
- Are these frictions likely to be quantitatively meaningful?

Production efficiency in reality

- The following paper provides an elegant empirical evaluation:
Hsieh and Klenow (2009), "Misallocation and Manufacturing TFP in China and India," Quarterly Journal of Economics
- Approach: use micro data to measure the dispersion of marginal products across firms within three countries (US, China and India) and to back out the gaps from the frontier
 - We have seen: in our theory marginal products are equalized everywhere
 - Their bottom line: *"When capital and labor are hypothetically reallocated to equalize marginal products to the extent observed in the United States, we calculate manufacturing TFP gains of 30%–50% in China and 40%–60% in India."*
 - To put it differently, even if the best China and India could do was to reduce their frictions to the level of the US (itself not perfect), they could raise their income levels by nearly 50%.

Production efficiency in reality

- So, in reality there are many reasons why GDP is not maximized conditional on the available amounts of capital and labor.
- Through the lens of our aggregate production function, this means that micro inefficiencies lead to lower “ A ”.
- Our approach: just *assume* a common $F(K, L)$, and see how far we can get. If we don't account for data well then A is likely to be important.

Outline

1. A neoclassical production model

The aggregate production function

Firm behavior

Household behavior

Equilibrium

2. Aside: A single good?

Planner's problem vs. competitive equilibrium

Empirical assessment

3. Using the production model to interpret cross-country income differences

Our production model and income differences

- Let's use our model to take a look at income differences.
- Suppose each country j has Cobb-Douglas production function:

$$Y^j = F(K^j, L^j) = A^j (K^j)^{1/3} (L^j)^{2/3}.$$

- We just discussed how—because of micro inefficiencies— A^j could differ across countries even if the true technologies are the same. But for now, suppose all countries have the same aggregate TFP:

$$A^j \equiv A \text{ for each } j.$$

- This is a useful starting point:
 - We will see how far we can go by attributing GDP differences to differences in K and L
 - Then ask (i) what is the residual to be accounted for by differences in TFP and (ii) where such differences themselves may originate from.

Theory meets measurement

- We can measure labor (for now population) $\bar{L}^j$ for each j .
- We can also measure $\bar{K}^j$ for each j . Briefly:
 1. Make “educated” initial guess K_0 .
Not important for our purposes how to exactly arrive at this.
 2. Get investment from national accounts, and depreciation from accounting conventions.
Propagate forward:

$$K_t = (1 - \delta)K_{t-1} + I_t = \sum_{s=1}^t (1 - \delta)^{t-s} I_s + (1 - \delta)^t K_0$$

Note: for $\delta \in (0, 1)$ and $T \rightarrow \infty$, this is independent of K_0 .

- Can we explain income differences across countries from their differences in $\bar{K}^j, \bar{L}^j$?

Production model and income differences

- Note that in intensive form, we have:

$$y^j = \frac{Y^j}{L^j} = A \left(\bar{k}^j \right)^{1/3}, \text{ where } \bar{k}^j = \frac{\bar{K}^j}{L^j}.$$

- We can plug in observed $\bar{k}^j$ and obtain a predicted output, $y^{j,predicted}$, which we can then compare with actual output, y^j .
- Let's normalize the US capital and output to 1, that is:

$$\bar{k}^{US} = 1 \text{ and } A^{US} \equiv A = 1.$$


- These imply $y^{US*} = 1$. All measurements are relative to the US.
- $\bar{k}^j = 0.8$ says country j has 0.8 times the US capital per person.
- What is the predicted output per person relative to the US?

k differences predict income differences

FIGURE 4.4

Predicted Per Capita GDP in the Production Model

Predicted per capita GDP, y^*




k differences generate income differences

- Burundi has much less capital per person than the US. Our model implies Burundi should also have less income.
 - But 1/100 of the US capital implies about 1/5 of the US income.
 - Intuitively, sharply diminishing returns to capital since $\alpha = 1/3$.
 - Put differently, capital is important, but not hugely important—and thus, raising capital can only increase output (per person) by so much.
- What is the actual y difference between Burundi and the US?

k differences undershoot income differences

FIGURE 4.5

The Model's Prediction for Per Capita GDP (U.S. = 1)



The model would have good fit if the dots were on the line. But we see that the predicted output is almost always above the actual output.

k differences undershoot income differences

- Why does our model overpredict output? A natural culprit is the equal-TFP assumption:

$$A^j = A^{US} \equiv A.$$

- Burundi is probably less productive than the US

$$A^j < A^{US}.$$

- If this is the case, they would produce less than what the US could produce with the same amount of capital.

$$A^j (\bar{k}^j)^{1/3} < A^{US} (\bar{k}^j)^{1/3}.$$


- **Q:** How low does A^j need to be to improve the fit of the model?

The model suggests large TFP differences

FIGURE 4.6


The U.S. and Chinese Production Functions

Output per person, y



- The lower line chooses A^{China} to fit exactly $y^{China} = A^{China} \bar{k}^{China}$.
- The required level is $A^{China} = 0.327$ (where as before $A^{US} = 1$).

The model suggests large TFP differences



The same exercise for all countries implies large A^i differences.

Taking stock

- k differences seem important for income differences. But not hugely important due to diminishing returns, $\alpha = 1/3$. The model suggests A differences are also important.
- But both k and A are endogenous objects.
 - Where do these k differences come from?
 - What exactly is A , and where do A differences come from?
- These questions will set the stage for the next few lectures ...

Roadmap: the rest of the first part

- We will begin by trying to understand **k differences**:
 - Two model frameworks to understand capital accumulation:
 1. Solow growth model
 2. Neoclassical (Ramsey) growth model
 - Preview: k accumulation is not key for cross-sectional income differences [as seen today] nor can it support long-run, persistent growth
 - But: develop useful methodological tools + still learn interesting things
Most notably: how should an economy decide between consumption and savings?
- After that, we will work to understand **A differences**:
 - Endogenous growth models, institutions and other “fundamental” causes
 - Preview: it's how you use the factors— A —not the factors themselves that is key for cross-country differences & the source of long-run growth

MIT OpenCourseWare
<https://ocw.mit.edu/>

14.05 Intermediate Macroeconomics
Fall 2025

For information about citing these materials or our Terms of Use, visit: <https://ocw.mit.edu/terms>.